arXiv:1404.5324v1 [q-bio.QM] 21 Apr 2014

For principled model fitting in mathematical biology

Thomas House*

April 23, 2014

Introduction

The mathematical models used to capture features of complex, biological systems are typically non-linear [8, 9], meaning that there are no generally valid simple relationships between their outputs and the data that might be used to validate them. This invalidates the assumptions behind standard statistical methods such as linear regression, and often the methods used to parameterise biological models from data are *ad hoc*.

In this perspective, I will argue for an approach to model fitting in mathematical biology that incorporates modern statistical methodology without losing the insights gained through non-linear dynamic models, and will call such an approach *principled* model fitting. Principled model fitting therefore involves defining likelihoods of observing real data on the basis of models that capture key biological mechanisms.

While few would argue with these general principles, traditionally it has been considered necessary either to sacrifice likelihood-based model fitting, or biological realism, for pragmatic reasons. I argue, however, that while some level of pragmatism is always necessary, mathematical biologists should be increasingly able to adopt a principled rather than a pragmatic approach when fitting models to data.

Of course, the massive and continuing increases in computational power and data availability play a major role in enabling the possibility of fitting mechanistic models to data statistically, but mathematical results are at least as important in my opinion. First, there is increasing use of stochastic modelling in biology, and stronger analytical results continue to be proved about these models, including the ability to relate standard differential-equation

*Warwick Mathematics Institute, University of Warwick, Coventry, CV4 7AL, UK.

models to underlying stochastic processes [1, 3]. Secondly, there has been extensive development of computationally intensive inference algorithms that can in principle deal with arbitrary likelihood functions [6, 4].

In the next few sections, I will develop three points in favour of a principled approach to model fitting: (1) accurate parameter estimation; (2) uncertainty quantification; and (3) the role of mechanism. These will be illustrated by examples using the SIR epidemic model with transmission rate β and recovery rate γ, technical details of which are collected in the Appendix. The examples are designed to be simple enough to support the point argued, but still to exhibit features of real-world problems.

1 Accurate parameter estimation

The primary reason for use of principled model fitting is to obtain accurate estimates of parameters, typically to predict likely future behaviour of a biological system. Suppose, for example, we are able to make a full observation of an infectious disease epidemic, including the onset of and recovery from disease, as may be possible in small populations such as boarding schools [5]. The first example, simplifying such a realistic scenario, is the general stochastic epidemic (1) with parameters $N = 500$, $\beta = 3$ and $\gamma = 1$. Let us further consider two possible data sets: (I) observation of the full epidemic; (II) observation until prevalence of infection is 50 for the first time.

Running Markov chain Monte Carlo (MCMC) with improper priors, using the tractable likelihood (3) for this model, it is possible to obtain parameter point estimates and 95% CIs of (I) $\hat{\beta} = 3.02[2.75, 3.31]$, $\hat{\gamma} = 0.925[0.844, 1.01]$ for full data and (II) $\hat{\beta} = 2.77[2.20, 3.53]$, $\hat{\gamma} = 0.718[0.419, 1.06]$ for early data only. As one would hope, the true values sit within the 95% CIs. Also, note that while this is a Bayesian procedure, similar results could be obtained in a frequentist framework, for example by using bootstrapping and numerical optimisation.

Suppose instead, we fit the standard differential-equation SIR model (7) to the data through numerical minimisation of the sum of squared differences (8). In this case the parameter estimates are (I) $\hat{\beta} = 2.16$, $\hat{\gamma} = 0.687$ for full data and (II) $\hat{\beta} = 3.70$, $\hat{\gamma} = 1.48$. Figure 1(i) shows that neither of these fits is a good description of the data. Figure 1(ii), however, shows that the principled method applied to early data provides useful predictions for future behaviour of the epidemic. Crucially, these predictions quantify uncertainty due to finite data and finite population size, and this is the second argument I wish to make in favour of principled model fitting.

(i) (ii)

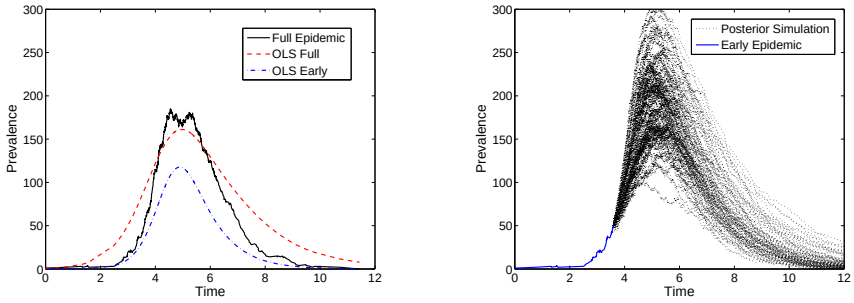

Figure 1: Results for fitting to a realisation of the general stochastic epidemic with $N = 500$, $\beta = 3$ and $\gamma = 1$, and τ defined as the first time at which $I = 50$. (i) Least-squares fitting to the full data (dashed line) and early data (dash-dotted line). (ii) 100 simulations from independent samples of the posterior obtained using Bayesian MCMC on the early data.

2 Uncertainty quantification

Another important feature of all real data is that it always leaves some questions unanswered, and quantifying the uncertainties that remain is almost always the second most important task after the strongest inferences that are possible have been made. Consider the case where an early epidemic in a large population is partially observed, so that there are a small number of time points at which prevalence is known. In this case, the 'gold standard' approach is generally considered to be MCMC with multiple imputation [10, 7], however this remains hard to implement and so here we use a tractable approximate likelihood (5) (thereby illustrating the important point that some level of pragmatism is often beneficial). Such data is often collected together with more selective sampling of individual cases to measure the course of infection. We simulate such a scenario using the model (1) with parameters $N = 10^5$, $\beta = 2$ and $\gamma = 1$, taking 11 samples of prevalence early in the epidemic as shown in Figure 2(i). We further sample the recovery times of $n = 150$ cases, with survivor function (1−the cumulative distribution function) shown in Figure 2(ii).

For these sources of data, there are linear relationships that can be used to obtain ordinary least-squares (OLS) parameter estimates. The first of these is that early in the epidemic, $\log(I(t)) = rt + \text{const.}$, where $r := \beta - \gamma$; the second is that the slope of the natural logarithm of the survival function for recovery is $-\gamma$. Figure 2(iii) shows that while these OLS methods yield sensible estimates of parameters, it is their confidence intervals that are misleading, since these suggest high confidence in a region of parameter space far from the true value. In contrast, the true value *is* in a region

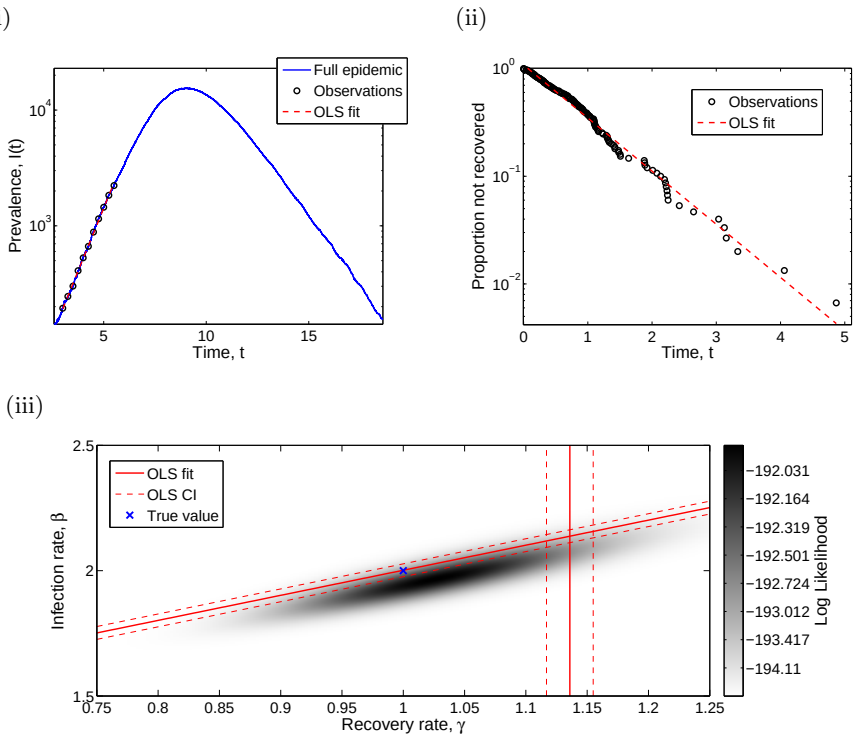

Figure 2: Results for fitting to (i) a realisation of the general stochastic epidemic with $N = 10^5$, $\beta = 2$ and $\gamma = 1$, where prevalence is measured at 11 time points at intervals $\delta = 0.5$ and (ii) $n = 150$ observations of times to recovery from infection. (iii) shows the joint likelihood density, the true parameter values and the estimates from OLS methods.

of high likelihood density, where the likelihood is formed from the product of (5) and the probabilities of observed recovery times.

3 The role of mechanism

The third point I wish to make is that even the most sophisticated statistical approach is neither a substitute for scientific understanding of the biological mechanisms behind data, nor an alternative to consideration of appropriate applied questions relating to prediction or intervention. Suppose we are in a situation where the only data available are three prevalence estimates early in the epidemic. In this case, we have some information about $r = \beta - \gamma$, but no independent information about γ. Figure 3(i) shows the behaviour of SIR models with different values of γ, as well as a pure birth process (i.e. $I \to I + 1$ at rate r with no recovery and no upper limit to the magnitude of

(i) (ii)

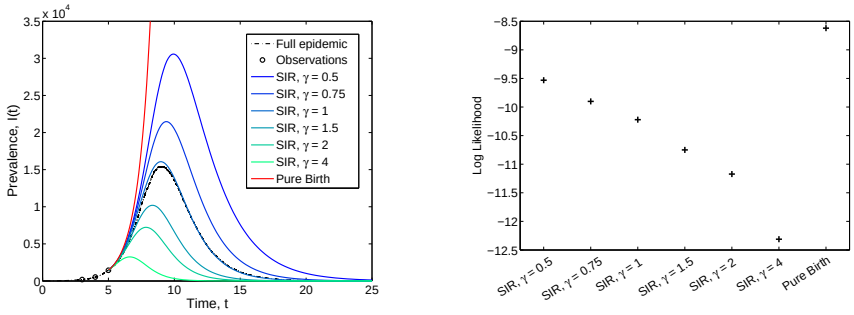

Figure 3: Results for fitting to a realisation of the general stochastic epidemic with $N = 10^5$, $\beta = 2$ and $\gamma = 1$, where prevalence is measured at 11 time points at intervals $\delta = 0.5$. (i) Predictions for fitted models with various values of γ and a pure birth process, and (ii) comparison of the likelihoods of fitted models.

I). (ii) shows that model likelihood is maximised for the pure birth process, but scientifically this is not the correct model and its long-term predictions (unlimited, never-ending growth of infection) are very far from reality.

I argue that therefore, in data-poor scenarios, the role of modelling remains to capture key mechanisms and provide predictions that are conditional on the values of unknown parameters. In fact, the complexity of biological systems means that typically data will always be inadequate to estimate absolutely every parameter of interest, meaning that mathematical biologists should generally be less keen than pure statisticians are to wield Occam's razor.

Conclusions

In conclusion, I have argued in favour of principled methods for model fitting in mathematical biology, which involve both statistical methods such as likelihood functions, and also the mechanistic models and insights of traditional mathematical biology. While biological systems will never be as predictable or amenable to precision measurement as physical ones, the development of models that are both predictive (together with appropriate quantification of uncertainty) and adequate for description of existing observations is now increasingly realisable. My perspective is that this trend will continue, and will be an enormously positive development for the field of mathematical biology.

References

[1] H. Andersson and T. Britton. *Stochastic Epidemic Models and Their Statistical Analysis*, volume 151 of *Springer Lectures Notes in Statistics*. Springer, Berlin, 2000.

[2] N. T. J. Bailey. *The Mathematical Theory of Epidemics*. Griffin, London, 1957.

[3] A. J. Black and A. J. McKane. Stochastic formulation of ecological models and their applications. *Trends in Ecology & Evolution*, 27(6):337–345, 2012.

[4] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors. *Handbook of Markov Chain Monte Carlo*. CRC Press, 2011.

[5] Communicable Disease Surveillance Centre (Public Health Laboratory Service). Influenza in a boarding school. *BMJ*, 1(6112):586–590, 3 1978.

[6] W. R. Gilks, S. Richardson, and D. J. Spiegelhalter. *Markov Chain Monte Carlo in Practice*. Chapman and Hall/CRC, 1995.

[7] C. P. Jewell, T. Kypraios, P. Neal, and G. O. Roberts. Bayesian analysis for emerging infectious diseases. *Bayesian Analysis*, 4(4):465–496, 2009.

[8] J. D. Murray. *Mathematical Biology I*. Springer, 3rd edition, 2002.

[9] J. D. Murray. *Mathematical Biology II*. Springer, 3rd edition, 2003.

[10] P. D. O'Neill and G. O. Roberts. Bayesian inference for partially observed stochastic epidemics. *Journal of the Royal Statistical Society A*, 162:121–129, Jan 1999.

A Technical Appendix

A.1 Individual-based model and likelihood

The underlying stochastic process is the *general stochastic epidemic* model [2], which consists of two integer-valued non-independent random variables in continuous time $S(t)$ and $I(t)$, such that $S(t) + I(t) \leq N$, where the integer N is the population size. This model has two real-valued parameters, λ and γ, which have dimensions of inverse time and determine the rates of processes of the Markov chain:

$$(S, I) \to (S - 1, I + 1) \text{ at rate } \lambda SI \ , \quad (S, I) \to (S, I - 1) \text{ at rate } \gamma I \ . \quad (1)$$

And so if the current state is (S, I) then the probability densities for the next event being an infection or recovery after time t are respectively

$$\rho_1(t) = \lambda SI\mathrm{e}^{-(\lambda S+\gamma)It} \ , \quad \rho_2(t) = \gamma I\mathrm{e}^{-(\lambda S+\gamma)It} \ . \tag{2}$$

We will consider the case where observations are a set of times T and events $\{e(t) \mid t \in T \ \& \ e(t) \in \{1, 2\}\}$, so that a likelihood function can be defined as

$$\mathcal{L}(\beta, \gamma) = \prod_{t\in T} \rho_{e(t)}(t) \ . \tag{3}$$

A.2 Early diffusion limit

For a population with large size N, the early epidemic prevalence $I(t) \ll N$ converges $I(t) \to Y(t)$, where $Y(t)$ obeys the stochastic differential equation

$$\frac{\mathrm{d}Y}{\mathrm{d}t} = (\beta - \gamma)Y + \left(\beta^2 + \gamma^2\right)^{1/2} Y\xi \ , \tag{4}$$

where $\beta := N\lambda$. If we make a series of observations $\{y_m\}$ of infectious prevalence at times $\{t_m\}$, then we can approximate the likelihood using a Gaussian process:

$$\begin{gathered}
\mathcal{L}(\beta, \gamma) = \prod_m \mathbb{P}[Y(t_{m+1}) = y_{m+1} | Y(t_m) = y_m] \ , \\
\mathbb{P}[Y(t+\delta) = y' | Y(t) = y] \approx \mathcal{N}\left(y' \middle| \mu = y\mathrm{e}^{(\beta-\gamma)\delta}, \ \sigma = \left(\beta^2 + \gamma^2\right)^{1/2} \mu\delta\right) \ , \\
\mathcal{N}(x|\mu, \sigma) := \frac{1}{(2\pi)^{1/2}\sigma}\mathrm{e}^{-(x-\mu)^2/(2\sigma^2)} \ .
\end{gathered} \tag{5}$$

A.3 Deterministic limit

In the limit of large N (or more strictly $I(t) \gg 1$) the stochastic process (1) converges on the well-known SIR equations

$$\frac{\mathrm{d}s}{\mathrm{d}t} = -\beta si \ , \quad \frac{\mathrm{d}i}{\mathrm{d}t} = \beta si - \gamma i \ , \tag{6}$$

where

$$s(t) := \frac{1}{N}\mathbb{E}[S(t)] \ , \quad i(t) := \frac{1}{N}\mathbb{E}[I(t)] \ . \tag{7}$$

If one approached data of the kind discussed in §A.1 with the equations (6), then an alternative to likelihood-based fitting would be to employ a 'least squares' approach, and choose parameters to minimise the function

$$\mathcal{E}(\beta, \gamma, i_0) = \sum_{t\in T} \left(I(t) - i(t; \beta, \gamma, i_0)\right)^2 \ . \tag{8}$$